MINDFORGE

AI Risk Management: Executive Handbook

November 2025

MINDFORGE

The MindForge Consortium

MAS Monetary Authority of Singapore

AIA HEALTHIER, LONGER, BETTER LIVES · BlackRock · citi · DBS · eastspring investments A Prudential plc company

Great Eastern An OCBC Company · GIC · GxS · HSBC · HSBC Life

income made yours · Julius Bär · Manulife · Maybank · MSIG

MUFG · Munich RE · OCBC · PRUDENTIAL · SMBC

standard chartered · STATE STREET · UBS · UOB

aws · Google Cloud · Microsoft

Supported by: accenture

With participation from:

abs The Association of Banks in Singapore · General Insurance Association · imas Investment Management Association of Singapore · Life Insurance Association Singapore Life is worth protecting. Invest in it. · SFA Singapore FinTech Association

Introduction

Artificial Intelligence (AI) is one of the most impactful technologies to be adopted in the financial services industry in recent years. It presents a range of opportunities for efficiency and better experiences for employees and customers, but also presents a range of risks if not well-governed. Project MindForge, of which this Handbook is a part, was launched in 2023 as the continuation of a multi-year legacy of proactive industry collaboration to address the responsible use of AI with the long-term leadership and support of the Monetary Authority of Singapore (MAS).

MAS issued the 14 FEAT (Fairness, Ethics, Accountability, Transparency) Principles for responsible AI use in the financial services industry in 2018.[1] Following the introduction of the FEAT principles, the Veritas Initiative was established by MAS and a consortium of financial institutions (FIs), consultancies, and technology companies to operationalise those principles.[2] Between 2020 and 2023, the Veritas Initiative co-created guidance to help FIs evaluate real-world solutions using AI and data analytics against the FEAT Principles, producing a Methodology and Toolkit. This Methodology is now used by FIs to implement AI and data analytics solutions responsibly.

Following significant advancements in Generative AI (Gen AI) technology in late 2022, the industry felt it necessary to assess the risks of this new technology, examine the FEAT Principles and Veritas Methodology to confirm their applicability to these risks, and adapt them where needed. This was the focus of Phase 1 of Project MindForge, which had four key outcomes. The first was to produce the first financial industry-specific taxonomy of Gen AI risks. The second was to review the FEAT Principles and the Veritas Methodology against these risks, identifying areas where additional considerations may be required to address the unique characteristics of Gen AI. The third was to provide an overview, based on the state of the art at that time, of the architectural and infrastructure considerations around responsibly developing and deploying Gen AI. The fourth was to develop two practical Gen AI use cases on risk management and compliance that demonstrated the application of responsible principles to new Gen AI tools. Phase 1 concluded with the publication of a whitepaper on the emerging risks and opportunities of Gen AI for banks in May 2024.[3] FIs are using the MindForge whitepaper to adapt their AI governance and risk management approaches to the challenges of Gen AI.

The Association of Banks in Singapore (ABS) elaborated on the MindForge whitepaper to publish a Handbook on Generative AI Guardrails in Banking in May 2025. This Handbook focuses on a selection of risks highlighted in the MindForge whitepaper and proposed specific, tangible guardrails for addressing them. Its work was an important input in the development of this Handbook.

Phase 2 of Project MindForge formally kicked off in November 2024. Its mission is to enable and facilitate FIs, at different levels of AI maturity, to scale AI with trust by adopting and operationalising AI governance and risk management across the enterprise: supporting industry AI use that is rapid, but responsible.

This Handbook, the first part of Phase 2 of Project MindForge, draws on the FEAT Principles, the work of the Veritas Initiative, and the risks identified in Phase 1 of Project MindForge to create a comprehensive guide to AI governance and risk management for the industry. It extends the work of MindForge Phase 1 to include traditional AI, Gen AI, and more recent technologies such as Agentic AI. Harmonising the range of good practices across the ecosystem into one Handbook will help make AI governance and risk management

[1] Read the FEAT Principles at https://www.mas.gov.sg/publications/monographs-or-information-paper/2018/feat
[2] Read about the Veritas Initiative, and see its publications, at https://www.mas.gov.sg/schemes-and-initiatives/veritas
[3] Read the story of Phase 1 of Project MindForge, and the 2024 industry whitepaper, at https://www.mas.gov.sg/schemes-and-initiatives/project-mindforge

straightforward, systematic, and universal. This Handbook is intended to accompany and support the implementation of the proposed MAS Guidelines on Artificial Intelligence Risk Management.

Industry alignment on the Handbook began in advance of the public kick-off in September 2024, when the consortium's 24 primary members endorsed the Handbook's scope and structure. The consortium convened three Working Groups made up of the primary members, tech partners, and a consulting partner, which drafted the text of the Handbook between February and July 2025. The scope and draft text were also reviewed and supported by the members of financial industry associations in Singapore through four engagement sessions in February, June, August, and November 2025, and two rounds of open feedback in November-December 2024 and August-September 2025. Over 100 financial organisations outside the primary consortium were engaged through these activities. After taking that feedback into account, this Handbook was formally launched at the Singapore Fintech Festival in November 2025.

The Handbook consists of three documents, of which this is the first. These three documents are:

AI Risk Management Executive Handbook (This Document). This document provides Considerations and Implementation Practices for governing AI across each Section in the Handbook's scope. It is intended as a resource for executives in the financial services industry.

AI Risk Management Operationalisation Handbook. This document provides detailed guidance on the operationalisation of each of the Implementation Practices recommended under each of the Handbook's Consideration. It includes illustrations of good practices from primary members, appendices, and other supporting materials.

AI Risk Management Handbook Implementation Examples. This document provides detailed case studies on individual financial institutions' experiences implementing AI governance and risk management.

These three documents are meant to be used in conjunction, and together make up the MindForge AI Risk Management Handbook.

MindForge is founded on a commitment to using AI responsibly, in a manner that manages its risks while leveraging its benefits. Governance and adoption are not, as concepts, in tension; in fact, widespread, rapid, and useful innovation in AI requires robust risk management and good governance. FIs that responsibly manage the risks of AI will be able to transform their businesses with the confidence that new technologies will behave as intended, follow the law, be secure, and protect their (and their employees and customers') data. It accelerates innovation and supports value realisation through measures that support observability, controllability, and oversight. It allows customers, employees, stakeholders, and society to trust FIs that use AI because they are confident that the technology's use will be fair, ethical, accountable, and transparent. Rather than impeding AI innovation, the practices outlined in this Handbook will enable it.

Acknowledgements

We would like to extend our thanks and recognition to the participants in the development of this Handbook.

Project MindForge is a collaborative industry initiative led by the Monetary Authority of Singapore (MAS) and delivered cooperatively by a consortium of FIs across the banking, insurance and capital market sectors, supported by consulting and technology partners and industry associations. Several members volunteered to take on additional responsibilities as leads and co-leads of the MindForge consortium's three working groups, which wrote the text of the Handbook.

The primary members of the MindForge consortium are (in alphabetical order):

- AIA
- BlackRock (Co-Lead for the "Data & AI" and "Enterprise" Working Groups)
- Citi
- DBS (Lead for the "Data & AI" Working Group)
- Eastspring Investments
- GIC
- Great Eastern Life
- GXS Bank
- HSBC
- HSBC Life
- Income Insurance (Co-Lead for the "Data & AI" Working Group)
- Julius Baer (Co-Lead for the "Data & AI" Working Group)
- Manulife
- Maybank
- MSIG
- MUFG Bank
- Munich Re
- OCBC
- Prudential (Co-Lead for the "Risk & Compliance" Working Group)
- SMBC
- Standard Chartered Bank (Lead for the "Enterprise" Working Group)
- State Street
- UBS
- UOB (Lead for the "Risk & Compliance" Working Group)

The development of the handbook was supported by the consortium's consulting partner:

- Accenture

The consortium was advised by its technology partners, which are:

- Amazon Web Services
- Google Cloud
- Microsoft

The consortium would like to acknowledge the contributions made by the approximately 200 individuals from the aforementioned institutions who participated in the Handbook's drafting.

Several financial industry associations and their membership provided guidance, inputs, feedback, and support throughout the development process. The five associations that were members of the consortium are:

- Association of Banks in Singapore
- General Insurance Association of Singapore
- Investment Management Association of Singapore
- Life Insurance Association Singapore
- Singapore FinTech Association

The consortium would like to acknowledge the contributions of these participating organisations and the numerous peers in the industry who contributed to this Handbook's successful development and refinement.

MindForge AI Risk Management and Governance Framework

Each Section in this Handbook corresponds to a component in the framework below, which represents a logical model to support FIs in implementing the Handbook.

Figure 0.0.1: MindForge AI Risk Management and Governance Framework

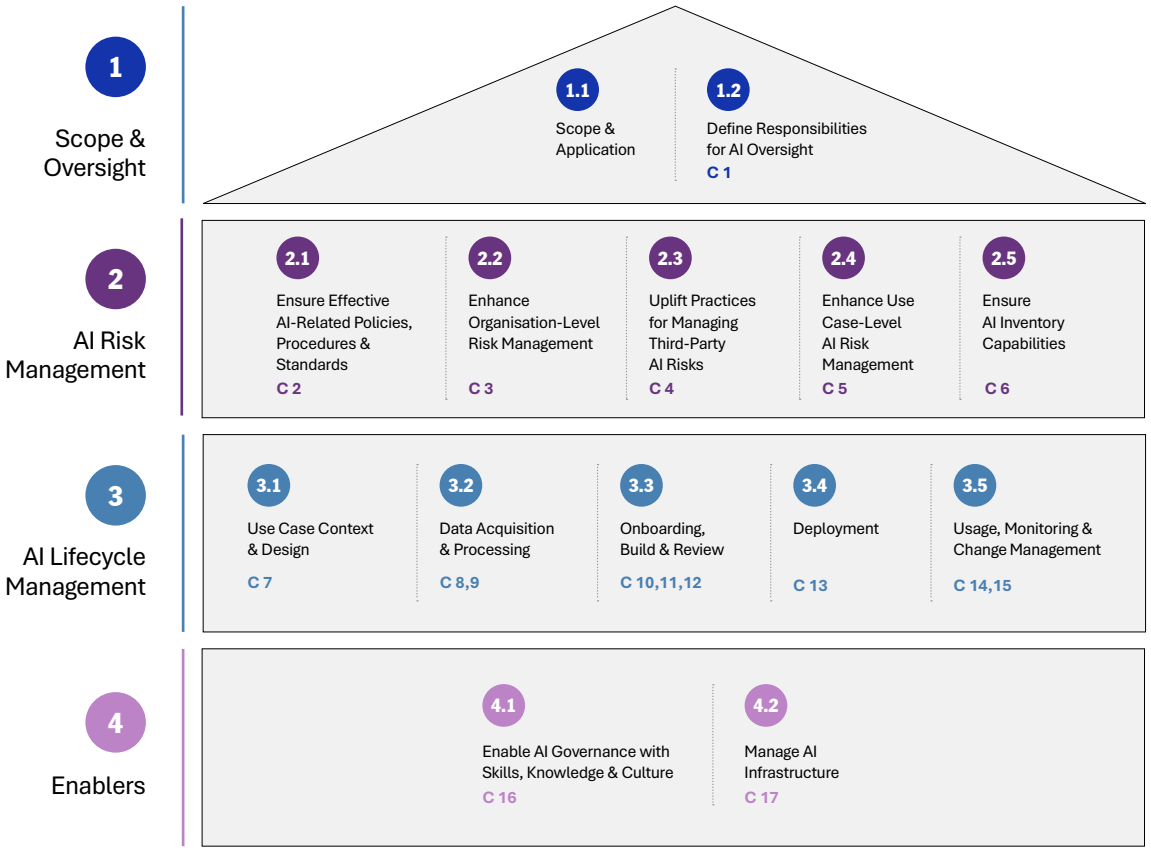

C – Considerations included in that Subsection

Section 1, **Scope and Oversight**, discusses some of the foundational concepts underpinning this Handbook. It also discusses how AI is overseen in an FI's operating model.

Section 2, **AI Risk Management**, discusses how FIs can measure, monitor, and mitigate the risks of AI by establishing policies, procedures, processes, and systems in their organisation.

Section 3, **AI Lifecycle Management**, discusses the key activities and considerations that can be applied to manage risks at each stage of the lifecycle of an individual AI use case.

Section 4, **Enablers**, discusses the foundational capabilities that can support AI risk management.

1. Scope and Oversight

1.1 Scope and Application

Organisations in Scope

This Handbook is addressed to all Financial Institutions (FIs) that use AI in their businesses.

This Handbook is addressed to FIs of all sizes. Although it makes specific references to Singapore's context, it is also intended to be globally relevant. As such, it aims to make recommendations that are aligned to global good practices beyond Singapore.

The Handbook is intended to be relevant to all functions within those organisations.

AI-Specific Risks in Scope

This Handbook defines practices for addressing AI-specific risks – new or enhanced risks arising from AI use that go beyond those of traditional software use in the context of an FI. The overall approach of this Handbook is to describe AI governance and risk management that is risk-based and proportionate, and that continuously improves as lessons are learned and as risks evolve. This Handbook addresses the risks posed to FIs by the use of AI in their businesses or by their vendors and service providers, and does not consider inbound risks caused by the use of AI by external parties.

This Handbook is designed to supplement and function in tandem with existing non-AI-specific risk management practices that FIs already have in place. These include practices for managing technology risk, cybersecurity, and risk management, many of which are governed by instruments listed in Appendix C of the Operationalisation Handbook. Where relevant, these non-AI-specific risks are referred to here, but are not replicated in this Handbook.

This Handbook aims to address governance and risk management for all types of AI, including traditional AI, Gen AI, and Agentic AI. In addition to the well-known risks of traditional AI, Gen AI and Agentic AI may introduce new or additional risks or challenges.

The Handbook took the work of MindForge Phase 1, which developed an AI risk taxonomy, as its starting point; an updated version of this taxonomy is provided in Appendix B of the Operationalisation Handbook. From this list of risks, the Association of Banks in Singapore (ABS) Handbook on Generative AI Guardrails in Banking highlighted ten "top risks", which were a particular focus:

- Unrepresentative or biased data inputs.
- Toxic and offensive outputs.
- Lack of AI risk awareness.
- Lack of use case, data and model governance.
- Inadequate human oversight.
- Inadequate feedback and recourse mechanisms.
- Hallucination/ Fabrication/ Confabulation.
- Overconfidence.
- Insufficient model accuracy/ soundness.
- Model degradation from unexpected use.

The landscape of AI risk is rapidly evolving, however, and in addition to the AI risk taxonomy in this Handbook, it is important that FIs and the industry overall can continue to consider new, emerging, or diminished AI-specific risks as they apply the Considerations in this Handbook.

Intended Audience

Within an FI, the Handbook is addressed specifically to:

- **Executives:** Decision-makers and leaders.
- **Builders:** Software developers, data engineers, data scientists, AI practitioners, systems integrators, and other technical specialists involved in the development, deployment, and use of AI.
- **Custodians:** Employees in oversight, governance, enablement, and risk management roles in an FI who apply AI governance and risk management policies and procedures and manage AI risks, either directly or in an enabling capacity such as talent, legal, or technology.
- **Use Case Owners:** Employees who are accountable for an AI use case.
- **Business Users:** Employees who use or apply AI use cases in the course of their business responsibilities.

Each FI may organise these functions differently, and under different titles, depending on its structure and needs (such as the concepts of the first, second, and third lines of defence). This handbook uses the above terms as generic terms of reference to common enterprise activities and aims to be relevant to all FIs irrespective of their internal organisation.

Structure of the Handbook

The Handbook is made up of 17 Considerations. These are the operative unit of the Handbook; each Consideration is a thematic recommendation that will support an FI in operationalising AI governance and risk management. Together, they form a checklist of actions that FIs can take to align with the approach in this Handbook. These Considerations are grouped thematically into several Subsections, which each begin with a brief overview that introduces its key concepts.

Under each Consideration are one or more Practices. These are specific actions that, when taken collectively and appropriately to the FI's context, can implement the Consideration. Each Practice is accompanied by a long-form text – the "Operationalisation Guidelines" – which describes that Practice in more detail and is omitted in the Executive Handbook.

Several Subsections are also enriched with an illustration contributed by a consortium member describing how the Practices in that Subsection are implemented in a real-world setting. These illustrations appear at the end of the Subsections where they are included.

Figure 1.1.1: Structure of Each Handbook Section

Define Responsibilities for AI Oversight

Consideration 1: Ensure that an AI governance operating model is clearly defined by leveraging and, as needed, uplifting the roles and capabilities of existing enterprise functions including the relevant roles from the Board, Senior Management, and operational governance, with sufficient operating effectiveness measures in place to support them.

Practice 1: Embed additional responsibilities for AI governance and risk management, as required, in relevant Board and Senior Management roles.

Operationalisation Guidelines:

Approach:

- Extend the roles and responsibilities of relevant Board members or bodies to include relevant AI-related actions, including the endorsement of key AI governance documents, ensuring that AI-specific skills are in place, and ensuring that AI risks are managed.
- Extend relevant existing Senior Management roles and responsibilities to include the implementation of effective AI governance and keeping the Board well-informed.

The roles of the Board and Senior Management are already well-defined in each FI; their core responsibilities for managing risk are described in Principle 11...

...

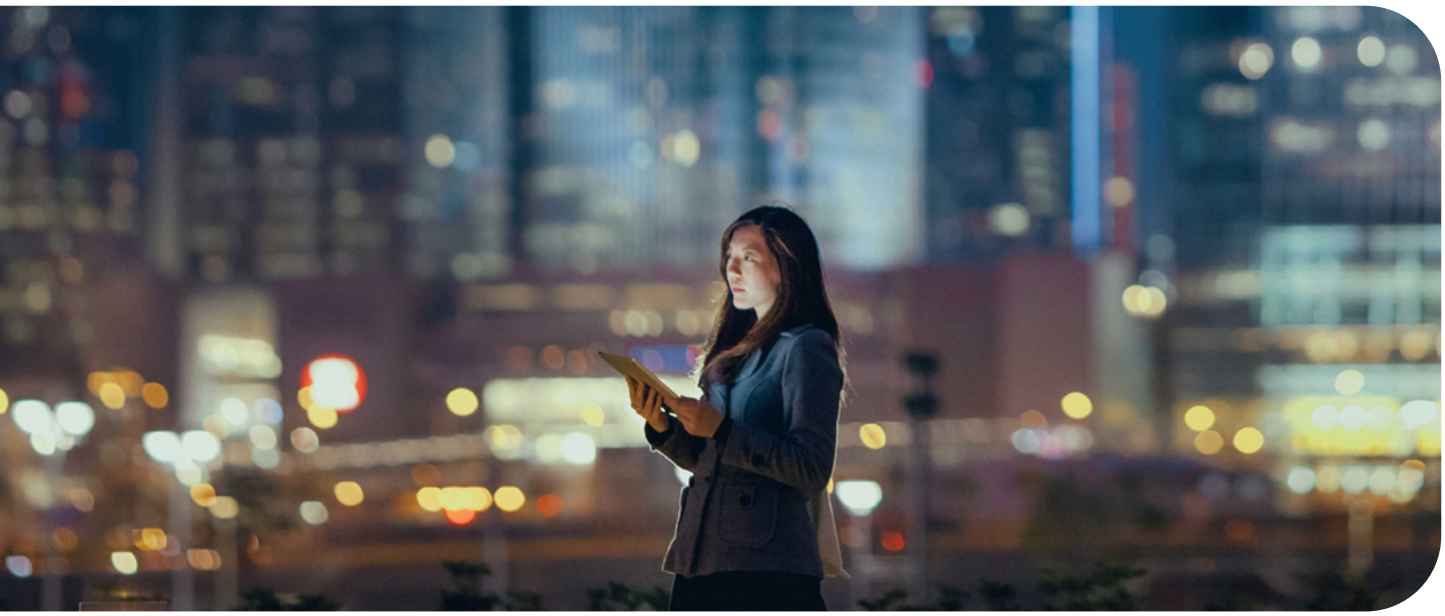

Relationship to Existing Enterprise Functions

AI is one of many technologies that FIs use, and FIs already have extensive enterprise functions in place for providing oversight, governance, and risk management beyond AI. Continuing to apply existing good practices in areas like risk management, data governance, procurement, software lifecycle management, talent, and cybersecurity is a key foundation for this Handbook – and for responsibly and effectively governing AI.

This Handbook was designed with the following goals in mind:

1. To build on, but not duplicate or replace, existing industry frameworks and good practices that apply to AI but are not AI-specific.
2. To describe only those considerations that are unique or additional to AI governance and risk management.

As a result, this Handbook does not describe activities or practices that are not specific or additional to the governance of AI. This should not be taken to imply that existing non-AI-specific practices are not also important to AI governance and risk management. FIs can continue to apply industry norms and good practices in risk management, data governance, procurement, software development, and cybersecurity, and can consult Appendix C of the Operationalisation Handbook for a list of other frameworks that were referenced in the development of this Handbook.

Relationship and Proportionality to Risk

This Handbook is based on the principle of proportionality between governance measures and AI risk. This risk-based approach underpins all of the Considerations in this Handbook and emphasises scaling risk management activities based on factors such as the nature of the FI's business, the scale and nature of its AI use, and its appetite for risk. This balanced approach encourages innovation and experimentation while focusing resources on managing the most significant risks.

The Handbook's Considerations are also based on the principle of relevance. Not all practices for AI governance and risk management will be feasible for all AI use cases, depending on the technologies used, the deployment pattern, or other situational factors. FIs can determine in context how to make this Handbook's recommendations relevant to their use of AI.

Relationship to FEAT and Veritas

The MindForge Handbook supports and builds upon the foundation of the FEAT Principles. These fourteen principles (logically grouped into Fairness, Ethics, Accountability, and Transparency) were issued in 2018 and continue to serve as a valuable reference to the industry. The overall direction set by FEAT remains an underlying philosophy underpinning the Considerations in this Handbook, and an indicative mapping to the individual principles of FEAT is provided in Appendix D of the Operationalisation Handbook. This Handbook is written to facilitate adherence to the FEAT Principles when its Considerations are applied.

The methodology developed by the Veritas Initiative is a detailed, widely accepted framework for implementing the FEAT Principles in practice. This methodology remains pertinent and effective in managing AI risk today; like FEAT, however, the fast-moving nature of the field means that practices have evolved substantially since it was written. Veritas, most notably, was not drafted with Gen AI or Agentic AI in mind and is focused on the governance of AI models, not AI use cases or the enterprise overall.

Major advances in Gen AI and Agentic AI since these frameworks were developed, however, have made it challenging to apply them as written to modern AI use cases. This Handbook is written with the intention of going beyond FEAT and Veritas to manage the risks of advanced AI, such as by extensively addressing the governance of the enterprise; viewing AI use cases holistically, rather than at the model level; and by considering a wider range of risks. The seven risk dimensions highlighted in Appendix B of the Operationalisation Handbook reflect this broader view.

Figure 1.1.2: Diagram of FEAT, Veritas, and MindForge

FEAT Principles

Foundational guidance that underpins Veritas and MindForge (2018)

Veritas TRUST IN AI

Veritas Initiative

Detailed methodology for implementing FEAT (2020-23)

MINDFORGE

Project MindForge

Updated and expanded approach across the enterprise (2023-)

Relationship to MAS Guidelines on AI Risk Management

This Handbook is intended to accompany and support the implementation of the proposed MAS Guidelines on Artificial Intelligence Risk Management. A mapping of this Handbook's Considerations to the proposed Guidelines, and a non-exhaustive list of other MAS frameworks that may be relevant in applying this Handbook, will be provided in Appendices D and C of the Operationalisation Handbook.

Relationship to Other Regulations

Leading global frameworks played an important role in the development of this Handbook, and are enumerated in Appendix C of the Operationalisation Handbook. The practices in this Handbook are closely aligned to global AI governance and risk management norms, and the Handbook overall can serve as a useful repository of industry leading practices that FIs around the world can consider. Adopting the Considerations in this Handbook can also support, but may not be sufficient to address, compliance requirements in other jurisdictions where AI is highly regulated. FIs may refer to existing data protection, risk management, and market conduct rules for non-AI-specific obligations that will continue to apply when using AI.

1.2. Define Responsibilities for AI Oversight

Effective oversight is a key element of AI risk management and governance, and serves to both enable and oversee the implementation of the other sections in this Handbook. Responsibility for AI oversight is typically integrated with existing governance functions, where it forms part of the FI's overall operating model. A key element of AI oversight is the role of the Board and Senior Management, who together take overall accountability for the FI's use of AI.

While some FIs may elect to hire dedicated AI risk management staff or establish new AI-specific oversight bodies or forums, none of the Considerations in this Handbook specify that FIs establish any new positions or create any new units.

Figure 1.2.1: Stylised View of the Enterprise and Related Handbook Subsections

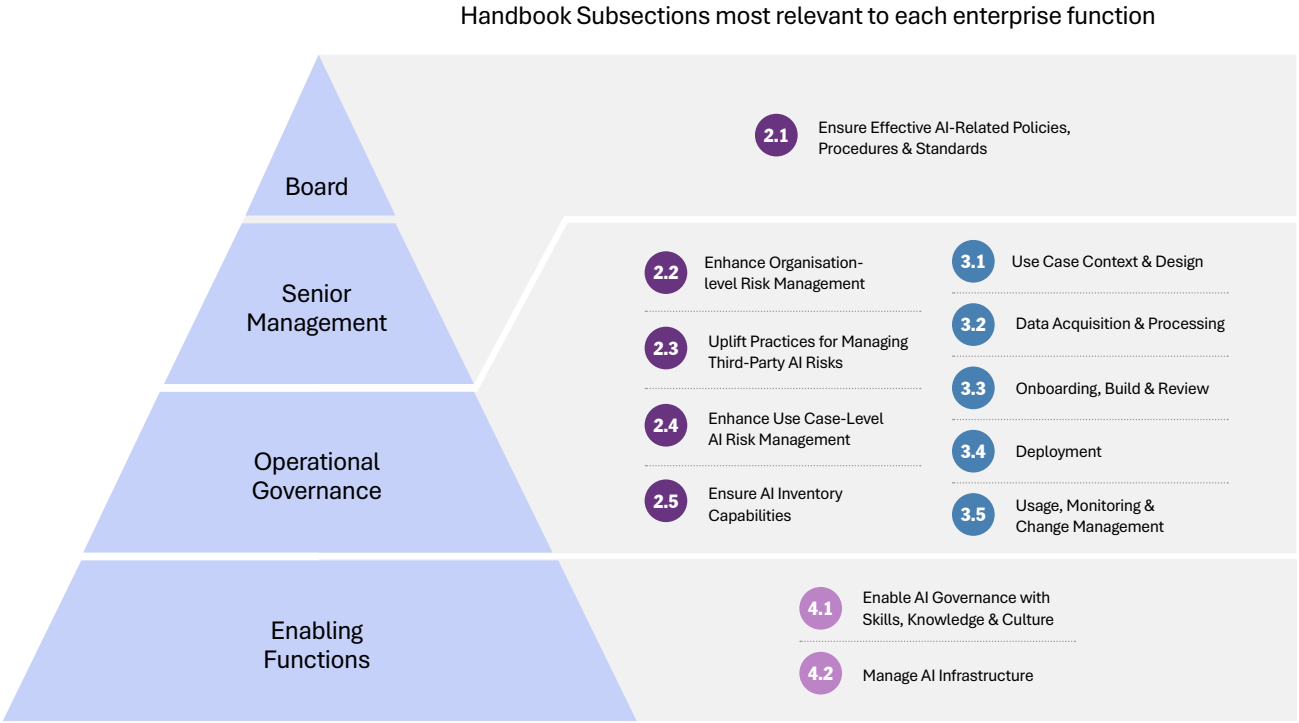

Consideration 1

Ensure that an AI governance and risk management operating model is clearly defined by leveraging and, as needed, uplifting the roles and capabilities of existing enterprise functions including relevant roles from the Board, Senior Management, and operational governance, with sufficient operating effectiveness measures in place to support them.

Practice 1: Embed additional responsibilities for AI governance and risk management, as required, in relevant Board and Senior Management roles.

Practice 2: Ensure that operational governance functions have clear roles and responsibilities assigned to operationalise AI governance and risk management activities across the enterprise.

Practice 3: Ensure that existing governance processes, forums, assets, and tools are updated to effectively enable AI governance and risk management.

Practice 4: Ensure that sufficient operating effectiveness and horizon-scanning measures are in place to monitor and improve the AI governance and risk management operating model over time.

2. AI Risk Management

2.1 Ensure Effective AI-Related Policies, Procedures, and Standards

FIs have a range of policies, procedures, and standards (collectively referred to here as "governance documents") in place in the organisation that may impact AI governance and risk management. These governance documents operationalise this Handbook's recommendations and establish a clear risk-based approach to managing AI throughout its lifecycle. Each FI will choose how best to update or supplement these governance documents to ensure that AI governance and risk management are effective, and may find, where they are already sufficiently mature, that further changes are not required.

Consideration 2

Ensure that governance documents define key AI-related concepts, processes, and responsibilities, and that they remain up-to-date and effective in supporting all aspects of the FI's approach to AI governance and risk management.

Practice 1: Ensure robust conceptual foundations for AI governance and risk management by establishing AI principles, defining key AI-related concepts, establishing frameworks for effective AI identification, and continuously improving these foundations over time as necessary.

Practice 2: Ensure that all aspects of AI governance and risk management are effectively institutionalised throughout the FI's governance documents, and that a process is in place to periodically review and reassess them.

2.2 Enhance Organisation-Level Risk Management

FIs that use AI in their business may encounter new or enhanced risks to the enterprise. FIs already have practices in place to define their organisational risk appetite and then identify, assess, treat, and monitor enterprise-level risks. FIs can leverage these existing practices to determine how best to address the risks of AI use and to update their own risk management approaches in a manner suitable to their needs and reflective of the broad variety of AI use cases in the organisation, each of which may have different control requirements.

Consideration 3

Enhance the organisational risk framework and risk appetite to include enterprise risks, strategies, and key risk indicators (KRIs) that track, monitor, and mitigate AI-specific risks.

Practice 1: Identify the new or enhanced risks of AI that are relevant to the enterprise and ensure that the enterprise risk taxonomy effectively captures them.

Practice 2: Assess existing enterprise risk controls for their fitness in addressing AI-specific enterprise risks, and uplift those controls where gaps exist.

Practice 3: Ensure that key risk indicators (KRIs) are in place to measure AI-specific risks and that relevant incidents, issues, or risk events are appropriately tracked and managed.

Practice 4: Ensure that effective monitoring is in place to identify AI-specific risk events or breaches of KRI thresholds to a degree proportionate to the FI's risk appetite.

2.3 Uplift Practices for Managing Third Party AI Risks

In addition to known, non-AI-specific risks inherent to the use of third-party technology, AI can present new or enhanced risks. These are related to the extent to which third party AI products and services may not be fully transparent, may evolve after deployment, may be challenging to assess for risk, may have incomplete information, or may lead to over-reliance. FIs can consider changes to their procurement, onboarding, and third-party risk management practices to address these new challenges.

To operationalise these practices, FIs can work with their vendors, can adopt industry-standard approaches where possible, and can be flexible to ensure that oversight is proportionate to risk.

Consideration 4

Uplift existing procurement and third-party risk management activities to address AI-specific risks, including disclosure templates, vendor assessment and procurement practices, change detection and notification, contracting practices, and ensure that teams have access to relevant expertise in AI.

Practice 1: Define, based on relevant AI-specific risks, a proportionate level of disclosure to seek from third party providers of AI products and services, and a process for assessing disclosures.

Practice 2: Ensure that processes and capabilities are in place for AI-specific risks to be evaluated at appropriate points in procurement, onboarding, and throughout the post-procurement lifecycle.

Practice 3: Identify new or modified AI components or features in third party products and services already introduced into the FI's technology ecosystem.

Practice 4: Consider whether contracts and licenses with third parties providing AI products and services are sufficient to clearly address AI-specific risks.

Practice 5: Ensure that teams with AI-specific legal, technical, and risk-management skills are involved in procurement, contracting, onboarding, or other third-party risk management activities as appropriate.

2.4 Enhance Use Case-Level AI Risk Management

To manage the risks of individual AI use cases, FIs can establish a framework for tiering their degree of risk materiality, applying controls in a risk-based approach, and reviewing these use cases for risk before and after deployment. As AI use cases are probabilistic, FIs can consider criteria like potential impact, complexity, and reliance to define AI risk materiality. Review can take a variety of forms depending on the maturity of the FI and the riskiness of the use case, ranging from peer review to a review from an individual or team elsewhere in the institution to review by an external party, at the FI's discretion. Doing so ensures that each AI use case can be governed proportionately.

Figure 2.4.1: Illustrative AI Risk Management Approach

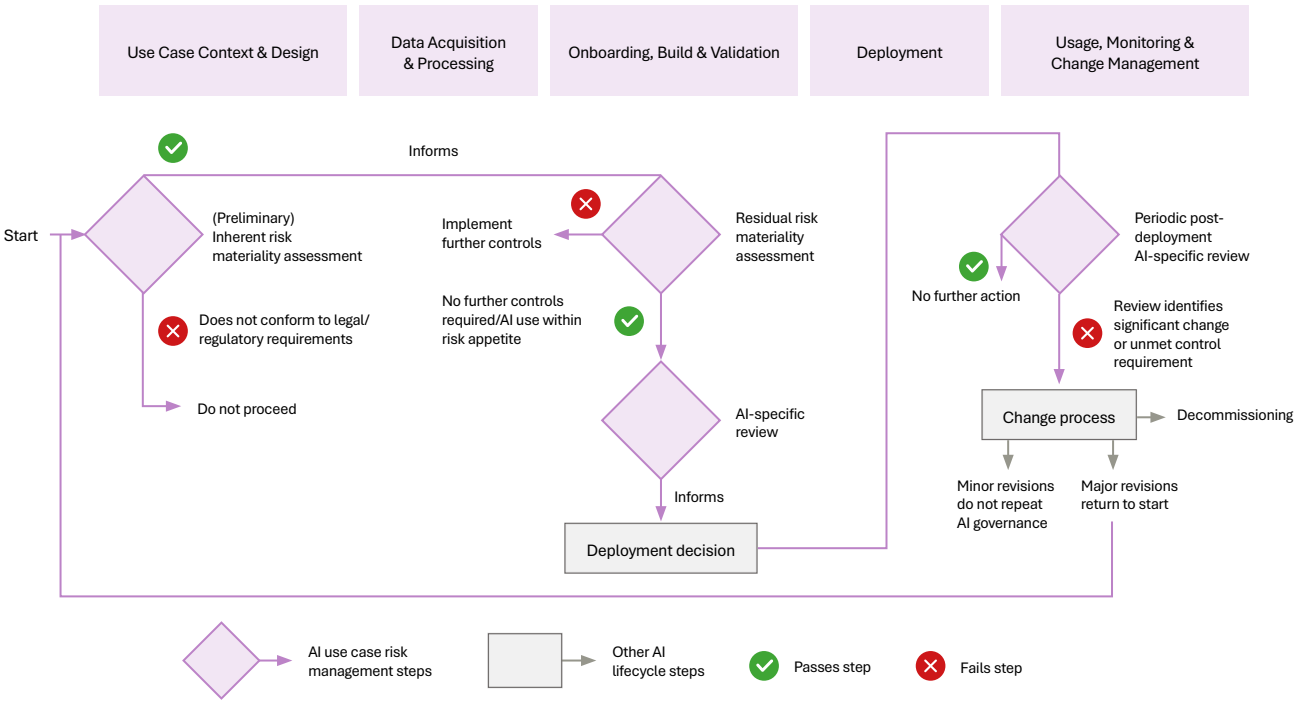

Consideration 5

Ensure that a framework is in place to manage the risks of each AI use case. This includes defining a risk materiality assessment approach, implementing a framework for inherent and residual AI risk assessments, applying controls that are commensurate with the risks identified, and conducting pre- and post-deployment AI-specific reviews as appropriate.

Practice 1: Define levels of risk materiality for AI use cases based on criteria relevant to the FI's context.

Practice 2: Define a process to assess the inherent risk materiality of AI use cases at the appropriate lifecycle stage, considering the fundamental characteristics of each use case.

Practice 3: Define a process to evaluate the residual risk materiality of AI use cases prior to deployment, considering the established controls and guardrails.

Practice 4: Identify, uplift, or create controls to be applied to each AI use case based on its risks and risk materiality.

Practice 5: Define an approach for conducting an AI-specific review of AI use cases prior to deployment, confirming the risks identified, the use case's risk materiality, and the appropriateness of risk mitigations.

Practice 6: Ensure that AI-specific reviews of AI use cases are conducted periodically post-deployment, with their frequency based on factors including the risk materiality of the AI use case.

2.5 Ensure AI Inventory Capabilities

An AI inventory is a repository of information on AI use cases. It ensures that AI use cases are used in the ways for which they were approved and supports overall, strategic risk management and governance of AI use in the FI.

Figure 2.5.1: Core Attributes of an Effective AI Inventory

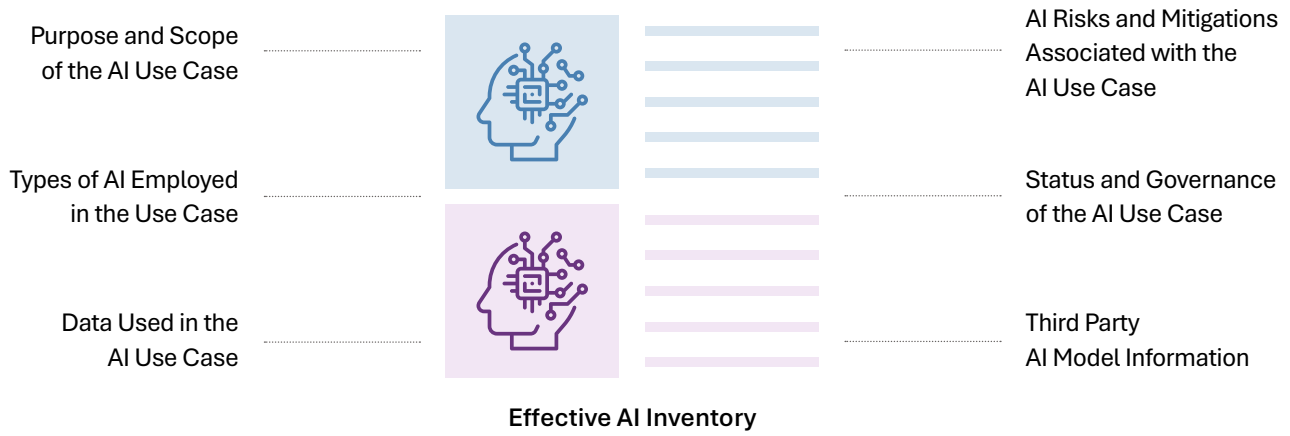

Consideration 6

Ensure that core AI-specific information on AI use cases is recorded in an inventory and ensure that a process is in place to maintain the AI inventory, so that information about new, updated, or decommissioned AI use cases is reflected accurately.

Practice 1: Ensure that a form of AI inventory, designed in consideration of existing inventory systems and practices to be suitable and proportionate for the FI's context, is in place to capture a core set of AI-specific information on AI use cases.

Practice 2: Ensure that processes are in place and that roles and responsibilities are defined such that the AI inventory is well-maintained and kept up to date.

3. AI Lifecycle Management

This Section provides guidance related to the lifecycle of individual AI use cases. It contains practices which may be relevant to some AI use cases and not to others, depending on their risk materiality, technology, and deployment pattern. Each FI can determine, on a per-use case basis, which practices are pertinent and proportionate to apply, and to what extent those practices can be applied.

Figure 3.0.1: AI Lifecycle

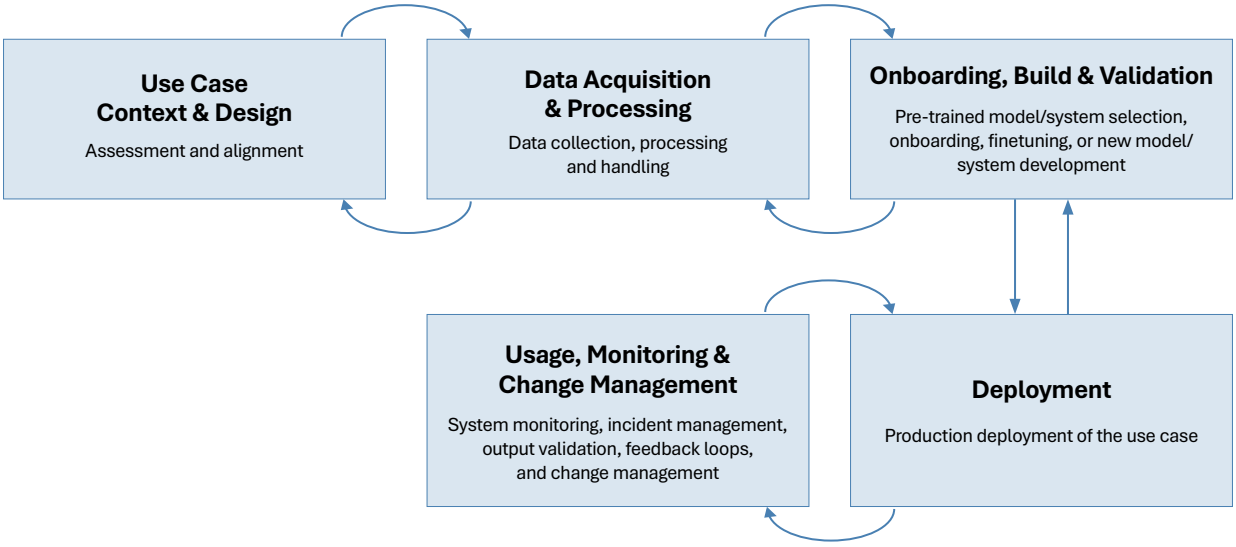

3.1 Use Case Context and Design

Use case context and design is the initial stage of the AI lifecycle where the plan for the use case is developed, is assessed for risk, and the FI ensures that it is aligned to their mission, values, priorities, and risk appetite. Many of these activities overlap with or can borrow from existing pre-development checks, and as such can leverage or build on existing use case governance.

Consideration 7

Assess the AI use case to ensure that the intended use is compatible with ethical, regulatory, and organisational standards, and determine the level of governance to be applied to the use case based on its inherent or expected risk materiality.

Practice 1: Establish ownership for the AI use case and ensure alignment with organisational standards and values for ethical and responsible AI use.

Practice 2: Perform an inherent risk materiality assessment to determine the risk tiering of the AI use case and to guide proportionate governance efforts.

Practice 3: Capture AI use case-related information in an AI inventory to enable transparency and support risk management.

Practice 4: Design the AI use case to operate with a proportionate and practical level of human oversight.

3.2 Data Acquisition and Processing

Data acquisition and processing is the stage where FIs procure and process the data that will enable the AI use case, and where risks introduced by the data itself or by its handling are mitigated.

Consideration 8

Evaluate whether the intended use of data in the AI use case is compatible with ethical, regulatory, and organisational standards.

Practice 1: Ensure that the use of data complies with ethical standards, regulatory requirements, and organisational policies or standards.

Practice 2: Ensure that the use of any third-party data complies with intellectual property rules, contractual obligations, and licensing rights.

Consideration 9

Adopt appropriate data management practices that address risks and limitations when processing data for AI use cases.

Practice 1: Ensure that data used for the AI use case is fit for purpose.

Practice 2: Justify the use of personal attributes in the AI use case.

Practice 3: Document metadata and data sources related to the AI use case in accordance with organisational data management policies and regulatory expectations.

Practice 4: Ensure that appropriate data access controls are implemented based on the nature of selected AI use case.

Practice 5: Establish clear ownership of any derived or transformed data to be used in the AI use case.

Practice 6: Identify and mitigate bias in training and test datasets.

3.3 Onboarding, Build, and Review

Onboarding, build, and review are three distinct stages; they may apply differently to third party and in-house developed AI use cases. Together, they involve developing and/or integrating AI components in the enterprise IT environment, and ensuring that they are in all cases effectively reviewed to ensure that they fall within the FI's risk appetite prior to deployment.

Figure 3.3.1: Illustrative Relationship Between Onboarding, Build, and Review

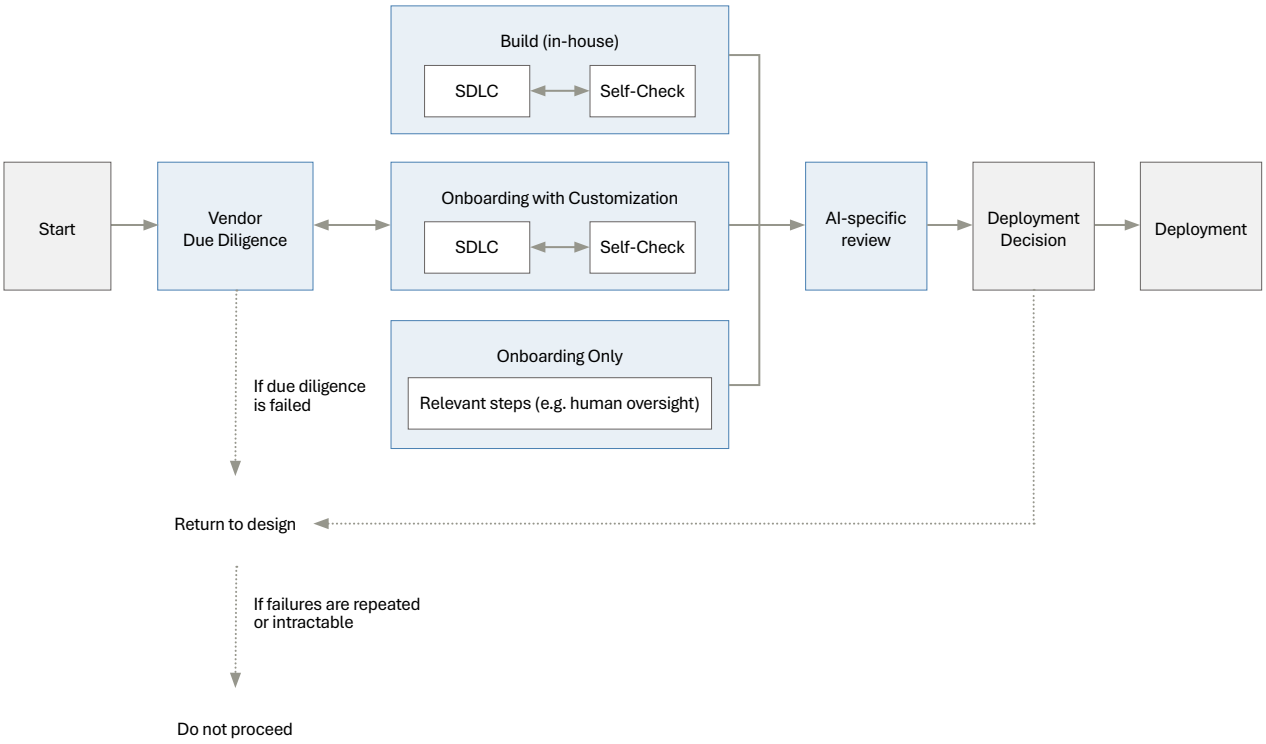

Consideration 10

Evaluate incremental AI-specific risks as part of the onboarding of third-party AI products and services within an AI use case.

> **Practice 1:** Conduct relevant use case-specific relevant due diligence during third-party AI onboarding, in line with organisational standards, to manage the risks of a third-party AI product or service.

Consideration 11

Ensure that the AI use case is built with appropriate guardrails and relevant metrics for effective performance and risk management.

Practice 1: Assess and select algorithms or features for the AI use case by considering its objectives and risks, including fairness, explainability, performance objectives, implementation complexity, and computational efficiency.

Practice 2: Identify and implement appropriate guardrails and controls during the development of AI use cases proportionately to the level and nature of the associated risks, to effectively manage and mitigate potential risks.

Practice 3: Define use case-specific risk-related metrics for assessing the AI use case for risks.

Practice 4: Evaluate and calibrate transparency measures based on the use case's risk materiality, degree of autonomy, and intended users, implementing proportionate design features and disclosures to support responsible and informed use.

Practice 5: Document key aspects of the AI build process, including data handling, model training and selection, and evaluation decisions to enable auditability and reproducibility.

Consideration 12

Conduct thorough testing and review prior to deployment to assess AI-specific risks and ensure that appropriate guardrails, controls, and governance have been observed.

Practice 1: Ensure that Builders conduct appropriate AI risk self-checks during development to test use case performance, verify the effectiveness of risk management activities, and identify and mitigate issues early in the development process.

Practice 2: Conduct an AI-specific review based on use case risk materiality prior to deployment to ensure that potential risks are identified and mitigated.

3.4 Deployment

Deployment is the stage at which the AI use case is put into a production environment. As they prepare to deploy a use case, FIs will complete key pre-deployment steps, including defining monitoring plans and ensuring effective security and integration.

Consideration 13

Develop monitoring and contingency plans for the use case prior to its deployment, and consider risk-informed deployment options.

Practice 1: In conjunction with other monitoring activities, ensure that a monitoring plan and safeguards/contingency measures are in place, along with the designation of an appropriate accountable person to address AI risks detected in monitoring.

Practice 2: Consider the need for a phased rollout to manage the AI use case's risks and progressively validate the use case's performance prior to full deployment.

Practice 3: Engage and equip users with targeted training and use case-specific resources to support responsible use and effective oversight.

Practice 4: Ensure that the AI use case is appropriately documented, that appropriate security and governance practices are applied, that relevant data retention is provided for, and that relevant approvals are obtained before deploying to production.

3.5 Usage, Monitoring, and Change Management

Usage, monitoring, and change management are three related lifecycle activities that continue after an AI use case is deployed. Usage involves ensuring that an AI use case is used in accordance with the FI's intentions. Monitoring involves checking that an AI use case's risks continue to be managed in an ongoing fashion, and change management is the practice of governing post-deployment modifications to ensure that risks are managed effectively.

Consideration 14

Conduct ongoing monitoring of the AI use case and its usage to ensure that it remains fit for purpose over time.

Practice 1: Periodically monitor and report on use case metrics related to AI risks, guardrail effectiveness, and changes in the use case's operating environment, as necessary and at a proportionate intensity and frequency, and address any issues identified.

Practice 2: Monitor and report on the quality, drift, and third-party risks associated with the use case's input and training data in an ongoing fashion, as necessary, after deployment.

Practice 3: Conduct periodic checks for changes to key aspects of the AI use case over time, including risk materiality, scope of usage, and key risks.

Practice 4: Conduct periodic AI-specific reviews after deployment to assess emerging post-deployment risks.

Practice 5: Ensure that the use case is operationalised with an appropriate degree of human oversight proportionate to its risk materiality or purpose.

Practice 6: Provide end users with avenues to enquire, give feedback, or request a review on AI decisions, where applicable, to support continuous improvement and build user trust.

Practice 7: Ensure that proportionate monitoring and analysis are in place to safeguard against security risks during system usage.

Consideration 15

Capture changes to AI use cases or their components to maintain traceability and ensure that changes with a material impact on risk are reviewed and approved through an effective change management process.

Practice 1: Establish AI change management process to ensure that changes to in-house or third-party use cases are appropriately tracked, reviewed, and approved before implementation.

4. Enablers

4.1 Enable AI Governance with Skills, Knowledge, and Culture

Appropriate skills, knowledge, and culture ensure that other measures on AI governance and risk management will be effectively and robustly implemented. FIs can leverage their existing measures around talent and conduct to ensure that these are fit for purpose for AI governance and risk management.

Consideration 16

Ensure that practices are in place to equip employees with the necessary AI governance and risk management skills, knowledge, and AI culture, while ensuring that teams involved in AI governance and risk management function are sufficiently representative.

Practice 1: Ensure that employees in relevant roles have the skills that they require to identify, mitigate, and track AI risks throughout the AI lifecycle.

Practice 2: Ensure that learning and literacy activities are sufficient to equip current and future employees with knowledge on AI capabilities, risks, and responsibilities appropriate to their roles in managing AI risk.

Practice 3: Ensure that practices, programmes, and policies related to culture and conduct are sufficient to foster a healthy AI culture around responsible, ethical, and safe AI use for current and future employees.

Practice 4: Ensure that AI governance and risk management activities involve a sufficiently representative and interdisciplinary group of employees who can effectively represent a range of perspectives on AI's risks and impacts.

4.2 Manage AI Infrastructure

AI use is enabled by a range of physical and digital infrastructure. FIs can continue to leverage existing good practices for technology risk to ensure that the risks of this infrastructure are well-managed.

Consideration 17

Support AI deployment by ensuring that supporting infrastructure is fit for purpose.

Practice 1: Ensure that the FI's AI-related infrastructure is suitable for managing scalability, availability, and security risks posed by the FI's use of AI.

Conclusion and Next Steps

This Handbook is an end-to-end guide to AI governance and risk management written by and for the financial services industry. Intended to be universally applicable – to FIs of all types and sizes, and to both well-established and emerging AI techniques – it provides a practical set of actions for mitigating AI-specific risks to people, businesses, and society, aligning AI to human values, and conforming to relevant laws and regulations. This foundation will support faster and better AI adoption across the financial services industry by creating user trust, supporting regulatory compliance, and facilitating more effective value realisation.

This Handbook contributes to the literature on AI governance and risk management by proposing useful definitions of core concepts, a risk-based, proportionate framework, and a flexible approach to process that is focused on uplifting practices that FIs already have in place. FIs are already using the recommendations in this Handbook to govern AI more effectively.

This Handbook is the first of several steps that the consortium will undertake to achieve its mission: to enable and facilitate FIs, at different levels of AI maturity, to scale AI with trust by adopting and operationalising AI governance and risk management across the enterprise, and supporting industry AI use that is rapid, but responsible. These next steps can include:

Ongoing Evolution

This Handbook is a living document; as AI rapidly evolves, these considerations will remain relevant only if they also continue to evolve in response to new regulations and technological developments. The consortium aims to continue to update the text of this Handbook, especially to ensure that it is aligned with the final version of the MAS Guidelines on Artificial Intelligence Risk Management.

The governance of emerging technologies like Agentic AI and other user-managed AI applications is a particular area of interest that the consortium will monitor further.

Training and Education

Effective AI governance and risk management involves building new skillsets – including training governance professionals to implement AI-specific frameworks like this Handbook. Building the talent pool for AI governance and risk management through effective training programmes will be a common good for the industry and is an area where the consortium can continue to collaborate.

Industry Toolkits

This Handbook highlighted the importance of consistent metrics and evaluation methods. Shared toolkits that implement generally acceptable metrics and methodologies can support the generalised adoption of this Handbook and can make its recommendations more implementable in practice.

Ecosystem Development

Institutionalising effective AI governance and risk management goes beyond adoption by FIs; it will involve the development of a robust ecosystem of fintech firms, technology providers, audit and assurance bodies, and other providers of ancillary capabilities that FIs will rely on in the long term.

The consortium remains committed to engaging with the AI governance and risk management ecosystem to support and promote innovation that can support its mission.